Robotics and Autonomous Systems 62 (2014) 675–684

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Identifying vegetation from laser data in structured outdoor environments

Kai M. Wurm *, Henrik Kretzschmar, Rainer Kümmerle, Cyrill Stachniss, Wolfram Burgard

University of Freiburg, Department of Computer Science, Georges-Köhler-Allee 79, 79110 Freiburg, Germany

A R T I C L E I N F O

Article history:
Available online 3 November 2012

Keywords:
Vegetation detection
Laser reflectivity
Self-supervised learning
Terrain classification
Autonomous navigation

A B S T R A C T

The ability to reliably detect vegetation is an important requirement for outdoor navigation with mobile robots as it enables the robot to navigate more efficiently and safely. In this paper, we present an approach to detect flat vegetation, such as grass, which cannot be identified using range measurements. This type of vegetation is typically found in structured outdoor environments such as parks or campus sites. Our approach classifies the terrain in the vicinity of the robot based on laser scans and makes use of the fact that plants exhibit specific reflection properties. It uses a support vector machine to learn a classifier for distinguishing vegetation from streets based on laser reflectivity, measured distance, and the incidence angle. In addition, it employs a vibration-based classifier to acquire training data in a self-supervised way and thus reduces manual work. Our approach has been evaluated extensively in real world experiments using several mobile robots. We furthermore evaluated it with different types of sensors and in the context of mapping, autonomous navigation, and exploration experiments. In addition, we compared it to an approach based on linear discriminant analysis. In our real world experiments, our approach yields a classification accuracy close to 100%.

1. Introduction

Autonomous outdoor navigation is an active research field in robotics. In most outdoor navigation scenarios such as autonomous wheelchairs, surveillance robots, or transportation vehicles, the classification of terrain plays an important role as most robots have been designed to drive on streets and paved paths rather than on surfaces covered by grass or vegetation. Failing to stay on paved roads introduces the risk of getting stuck and additionally increases wheel slippage, which may lead to errors in the odometry. Therefore, the ability to robustly detect vegetation is important for safe navigation in any of the above-mentioned situations.

In this paper, we propose a novel laser-based terrain classification approach that is especially suited for detecting low vegetation. Such low vegetation typically occurs in structured outdoor environments such as parks or campus sites. Our approach classifies the terrain based on laser scans of the robot's surroundings in order to allow the robot to take the classification result into account during trajectory planning. Low vegetation poses a challenge for laser-based terrain classification since the variance in the measured distances is often small and thus it is hard to detect it from range measurements alone. Therefore, our approach exploits an effect that is well known from satellite image analysis: Chlorophyll, which is found in living plants, strongly reflects near-IR light [1]. Mobile robots are often equipped with laser range finders such as the popular SICK LMS scanners. These devices emit near-IR light and return the range to the object they measure along with the reflectivity of the surface. Our work formulates the detection of terrain as a classification problem that uses reflectivity, incidence angles, and the measured distances as inputs to identify whether a measured surface corresponds to vegetation. In addition to that, we provide a way to gather training data in a self-supervised way and thus eliminate the cumbersome work of manually labeling training examples.

In this work, we evaluate two sensor setups: rotating laser scanners capturing 3D pointclouds and laser scanners mounted at a fixed angle. We show that classifiers can be trained in a self-supervised way using a vibration-based classification approach to label training data. To integrate classification results into a representation of the environment, we apply a probabilistic mapping method similar to occupancy grid mapping [2]. In our experiments, we illustrate that our approach can be used to accurately map vegetated areas and do so with a higher accuracy than standard techniques that are solely based on range values (see Fig. 1). We furthermore present applications to autonomous navigation in structured outdoor environments in which robots benefit from the knowledge about the vegetation in their surroundings.

This paper is organized as follows. After discussing related work, we will briefly describe support vector machines, which we use for classification. In Section 4, we present our approach to terrain classification using laser reflectivity. Section 6 describes the mapping approach used to integrate multiple measurements of

* Corresponding author.
E-mail address: wurm@informatik.uni-freiburg.de (K.M. Wurm).

0921-8890/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.robot.2012.10.003

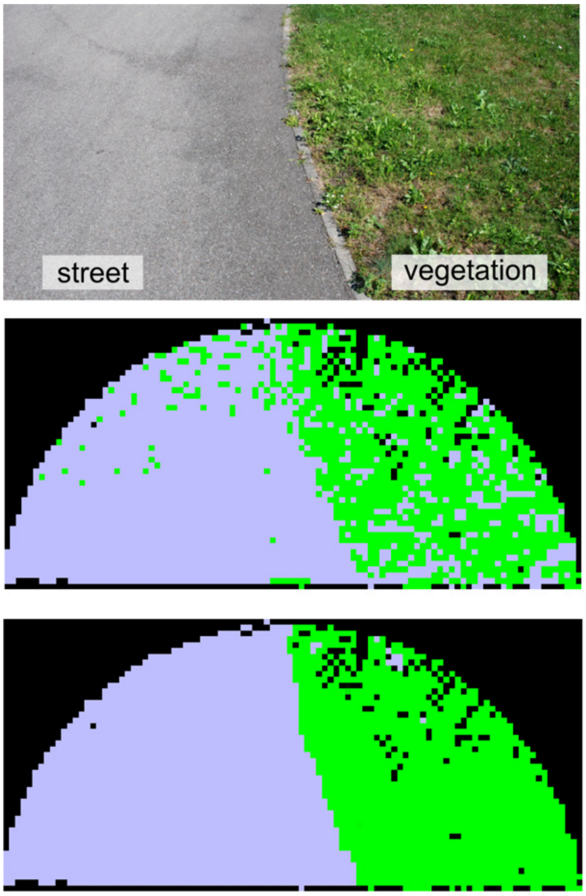

Fig. 1. A street and an area containing grass (top) and typical classification results obtained based on range differences (middle) and remission values (bottom). Shown is a bird's eye view of a 3D scan of the area depicted at the top with a maximum range of 5 m. Whereas points classified as street are depicted in blue, points corresponding to vegetation are colored in green. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the same region. Finally, in Section 9 we present the experimental results obtained with real data and with real robots navigating through our university campus and neighboring areas.

2. Related work

There exist several approaches for detecting vegetation using laser measurements. Wolf et al. [3] use hidden Markov models to classify scans from a tilted laser scanner into navigable (e.g., street) and non-navigable (e.g., grass) regions. The main feature for classification is the variance in height measurements relative to the robot height. Other approaches analyze the distribution of 3D endpoints in a sequence of scans [4–6]. These algorithms are able to detect various types of obstacles such as tree trunks, high grass, or bushes. However, flat vegetation, such as a freshly mowed lawn, cannot be reliably detected using this feature alone.

It seems intuitive to use color cameras to detect vegetation in the environment. Manduchi, for example, uses a combination of color and texture features to detect grass in camera images [7]. The main drawback of using cameras is that they are sensitive to lighting conditions. Shadows, for instance, can have a strong influence on the appearance of vegetation. A more robust classifier can be derived using the Normalized Difference Vegetation Index (NDVI). This value is based on the difference between red and near-infrared light reflectance and is a strong indicator for the presence of vegetation, see [1] for a discussion on this topic. On a mobile robot, the NDVI can be determined using a calibrated combination of a regular and an infrared camera [8] or a single camera that uses a filter mask to capture both frequency bands at different positions on the same sensor [9]. Unfortunately, such a sensor setup is rarely available on a mobile robot and, being a passive sensor, still depends on the presence of ambient light.

A combination of camera and laser measurements has been used to detect vegetation in several approaches [8,10–12]. In a combined system, Wellington et al. [12] use the remission value of a laser scanner in addition to density features and camera images as a classification feature. However, they do not model the dependency between remission, measured range, and incidence angle. Probably due to this fact, they found the feature to be only "moderately informative". In contrast to that, we show in our work that remission is highly informative if it is considered jointly with the measured range and the incidence angle of the individual laser beams.

The approach that is closest to our approach has been proposed by Bradley et al. [8]. Vegetation is detected using a combination of laser range measurements, regular and near-infrared cameras. Vegetation is recognized using the NDVI computed from both camera images. 3D laser measurements of the environment are projected into the camera images. A classifier is then trained using the vegetation feature and features from the distribution of 3D endpoints. According to the authors, the approach yields a classification accuracy of up to 95% but requires sophisticated camera equipment. In contrast to such systems that consider cameras in combination with laser range finders or infra-red cameras, our approach uses a laser scanner as the sole sensor. Our system is independent of lighting conditions and can be used on a variety of existing robot systems that are often already equipped with a laser range finder. We evaluated our approach in the context of detecting flat vegetation at ranges of up to 5 m. In this setting, we achieved classification results with an accuracy of over 99% in all our experiments.

Terrain types have also been classified using vibration sensors on a robot [13–16]. In these approaches, the robot traverses the terrain and the induced vibration is measured using accelerometers. The measurements are usually analyzed in the Fourier domain. Sadhukhan et al. [15] presented an approach based on neural networks. A similar approach is presented by DuPont et al. [14]. Brooks and Iagenemma [13] use a combination of principal component analysis and linear discriminant analysis to classify terrain. More recently, SVMs have been used by Weiss et al. [16,17]. Vibration-based approaches typically offer highly accurate classification results. The drawback of such methods, however, is that only the terrain the robot is moving on can be classified and not the terrain in front of the robot. Thus, this information cannot be integrated well into the path planning system of a mobile robot.

There exist previous approaches that apply self-supervised learning to classify terrain and detect obstacles. In the context of the LAGR-program (Learning Applied to Ground Robots), a number of methods were developed that exploit terrain knowledge in the surroundings of the robot to predict surface terrain in the far range. These near-to-far approaches use color information [18,19], 3D geometry information [20], or texture information [21,22]. Self-supervised learning was also used by Dahlkamp et al. [23] in a vision-based road detection system. Here, laser measurements are used to identify nearby traversable surfaces. This information is then used to label camera image patches in order to train a classifier that is able to predict traversability far away from the robot. In our approach, we adopt the idea of self-supervision to generate labeled training data. We apply a vibration-based classifier to label laser measurements recorded by the robot. This labeled dataset is then used to train a laser-based vegetation classifier. Both classifiers used in our approach have been implemented using support vector machines, which will be introduced in the following section.